

# Curve Trees: Practical and Transparent Zero-Knowledge Accumulators

Matteo Campanelli<sup>1</sup>      Mathias Hall-Andersen<sup>2</sup>

(preliminary version)

June 13, 2022

**Abstract.** In this work we propose a new accumulator construction and efficient ways to prove knowledge of some element in a set without leaking anything about the element. This problem arises in several applications including privacy-preserving distributed ledgers (e.g., Zcash) and anonymous credentials. Our approaches do not require a trusted setup and significantly improve on the efficiency state of the of the art.

We introduce new techniques inspired by commit-and-prove techniques and combine shallow Merkle trees, 2-cycles of elliptic curves to obtain constructions that are highly practical. Our basic construction—which we dub *Curve Trees*—is completely transparent (does not require a trusted setup) and is based on simple standard assumptions (DLOG and Random Oracle Model). It has small proofs and commitments and very efficient proving and verification time. Curve trees can be instantiated to be efficient in practice: the commitment to a set (accumulator) is 256 bits for any set size; for a set of size  $2^{32}$  a proof is approximately 2KB, a verifier runs in  $\approx 160$ ms (easily parallelizable to  $\approx 80$ ms) and a prover in  $\approx 3.6$ s on an ordinary laptop.

Using our construction as a building block we can construct a simple and concretely efficient anonymous cryptocurrency with full anonymity set. We estimate the verification time to be  $\approx 320$ ms (and trivially parallelizable to run in  $\approx 160$ ms) or  $< 10$ ms when batch-verifying multiple ( $> 100$ ) transactions simultaneously. Transaction sizes are  $< 3$ KB. Our timings are competitive with those of the approach in Zcash Sapling and trade slightly larger proofs (proofs in Zcash are 0.2KB) for a completely transparent setup.

## 1 Introduction

Zero-knowledge proofs are a cryptographic primitive that allows one to prove knowledge of a secret without revealing it. In many applications the focus is on proofs that are short and with efficient running time.

One of the rising applications of zero-knowledge is in *set-membership*: given a short digest to a set  $S$ , we want to later show knowledge of a member in the set without revealing the latter. This primitive is useful in domains such as privacy-preserving distributed ledgers, anonymous broadcast, financial identities and asset governance (see, e.g., discussion in [BCF<sup>+</sup>21]).

**Limitations of prior work.** Our focus in this work is on solutions that are highly *practical*. That is, solutions with short proving/verification time and short proofs. While efficient solutions to zero-knowledge set-membership already exist, we argue that they have limitations. In particular, either they still have a high computational/communication cost (we elaborate in Section 1.2 where we compare to transparent polynomial commitments and ring signatures [LRR<sup>+</sup>19]) or they rely on proof systems that are *non-transparent*. The latter means that, in order for the system to be bootstrapped, it is necessary to invoke a trusted authority. This is true for example in ZCash (Sapling) [HBHW21] and in [CFH<sup>+</sup>21]. While we can partly overcome this issue by emulating the trusted authority through a large-scale MPC, this is still highly expensive, both computationally and logistically<sup>1</sup>. Other solutions, such as [BCF<sup>+</sup>21, CHA21], mitigate this problem by requiring a trusted setup for parameters that are reusable in other cryptographic settings (an RSA modulus). This, however, still requires invoking a trusted authority or arranging a parameter-generation ceremony [CHI<sup>+</sup>20], which may not always be viable. We then turn to solutions that are fully transparent and still very efficient.

<sup>1</sup> <https://z.cash/technology/paramgen/>

**Our contributions.** Our main contribution is a concretely efficient construction for proving private set-membership with a fully transparent setup. Specifically we design a new data structure, CURVE TREES, that supports concretely small commitment to a set and where we can show set membership in zero-knowledge and with a small proof.

The design of a curve tree is simple and relies on discrete logarithm and the random oracle model (ROM) for its security. A curve tree can be described as a shallow Merkle tree where the leaves are points over an elliptic curve (and so are internal nodes). To hash, at each level we use an appropriately instantiated Pedersen hash alternating curves at each layer (we require a 2-cycle of curves). To prove membership in zero-knowledge we use commit-and-prove<sup>2</sup> capabilities of Bulletproofs and leverage the algebraic nature of our data structure. Our curves can be instantiated with existing ones in literature (see “Supported Curves” in Section 3.1). A curve tree has *structure-preserving* features[ACD<sup>+</sup>16] in that we never need to use any combinatorial hash (e.g., SHA) to convert representation of elements at each level or use their bit decomposition. While we focus on accumulators and set membership, our approach can straightforwardly be applied to opening of vectors rather than sets obtaining an “index-hiding” vector commitment.

Using our construction as a building block we can construct a simple and concretely efficient anonymous payment system with full anonymity set<sup>3</sup> and *transparent setup*. We dub this payment system  $\mathbb{V}\text{Cash}$ <sup>4</sup>. In  $\mathbb{V}\text{Cash}$ , the constraint system used for the zero-knowledge proof of a “spend” transaction is 20x smaller than that in ZCash Sapling (the currently deployed version of ZCash).

Concretely for two inputs/two outputs and anonymity sets of  $2^{32}$  (like in Zcash) our confidential transactions ( $\mathbb{V}\text{cash}$ ) require participants to compute/verify two Bulletproofs proofs of  $\approx 5000$  constraints each. Verifying each of the proofs in parallel (2 cores) in batches of at least 100 transactions (e.g. when verifying the validity of all transactions in a block) yields a very practical per-transaction verification time of  $< 10$  ms<sup>5</sup>. Transaction sizes are  $< 3$  KB. Our timings are competitive with those of the approach in ZCash Sapling and trade slightly larger proofs (proofs in ZCash are 0.2KB) for a completely transparent setup and simpler curve requirements.

## 1.1 Technical Overview

**Preliminaries: elliptic-curves and SNARK-native relations** In the following we assume that the reader is familiar with elliptic curves (see also Section 3.1 and notation in Section 2.1).

We informally say that a relation is “SNARK-native” to prove for a specific (SNARK) proof scheme if it can be “naturally represented in the constraint system” (a constraint systems is a representation of a relation we aim to prove). For example, we usually consider Pedersen hashing (and commitments) to be native to Bulletproofs (instantiated in the right curve). In fact we can prove we know the scalar representation  $\vec{u}$  of a group element  $U = \sum_i [u_i]G_i$  through roughly  $|\vec{u}|$  constraints<sup>6</sup>. Notice that for this operation to be actually native, each of the scalars  $[u_i]$  should be elements in the scalar field of the curve to which  $U$  and the  $G_i$ -s belong to.

**Starting point: shallow Merkle trees** As a warm up we will ignore zero-knowledge for most of this overview and then show how to account for it. Our starting point are *shallow* Merkle trees, i.e. Merkle trees with a general branching factor  $\ell \geq 2$ . Let us consider a balanced tree of depth  $d$ . This has  $N = \ell^d$  leaves (and it is encoding a set of an equal number of elements). We can represent this tree as labeled. Each internal node  $v'$  is labeled with the hash  $H(v_1, \dots, v_\ell)$  of the concatenation of its children; each leaf is labeled by its own value  $v$ . The root of this Merkle tree is public and represents the commitment to the set of elements. An internal node  $\text{rt}^{(i)}$  at level  $i$  can be seen as a root to a subtree branching from it. We denote by 0 the “lowest” level, to which the leaves belong, and by  $d$  the level to which the root belongs to (so we denote it by  $\text{rt}^{(d)}$ ).

One of the main advantages of trees with a high branching factor is that we may afford in practice a linear dependence on the depth. For a concrete vector size such as  $N = 2^{32}$ , we can choose  $\ell = \sqrt[4]{N} = 256$  and obtain a depth  $d = 4$ .

The straightforward approach to opening a leaf  $v^{(0)}$  in a Merkle tree opening provides the specific leaf together with the  $\text{rt}^{(i)}$ -s, the internal nodes along the path from  $v^{(0)}$  the root, and the sets  $\text{Siblings}(\text{rt}^{(i)})$  of siblings of each

<sup>2</sup> In the sense of the commit-and-prove building blocks in LegoSNARK [CFQ19] and in the work by Lipmaa [Lip16].

<sup>3</sup> An anonymity set can be seen as the subset of existing transactions a spent transaction can be narrowed down to. We say that a protocol supports a full anonymity set then it this set consists of the whole history of transactions so far.

<sup>4</sup> As a reference to both ZCash and Veksel [CHA21] from which it borrows part of its design.

<sup>5</sup> We use page 32 in [BBB<sup>+</sup>17] for our estimates. Our calculations are available at <https://gist.github.com/matteocam/377832d7b1f86cd0eac81149e9e65bdb>.

<sup>6</sup> While this is an informally defined notion, we contrast the Bulletproof example with the approach applying JubJub in Zcash Sapling. The latter is not *native* for Groth16 instantiated with BLS12-381 as it requires additional constraints for bit decompositions rather than directly describing the multiexponentiation.

$\text{rt}^{(i)}$ . This way, anybody can verify membership of the leaf by hashing the siblings at each level. The size of the opening certificate is roughly  $\ell \cdot d$ .

We would like to compress the communication complexity even further using SNARKs. However we are looking for better tradeoffs than what we can obtain by “plugging the whole tree opening inside the SNARK” (we discuss this more in the related work).

**Our basic blueprint** Our high-level solution stems from this insight. Instead of providing all siblings at opening time, we can just provide the internal nodes  $\text{rt}^{(1)}, \dots, \text{rt}^{(d)}$  together with short SNARKs  $\pi^{(1)}, \dots, \pi^{(d)}$ <sup>7</sup>. Each proof  $\pi^{(i)}$  should show “knowledge of the appropriate siblings”, namely of  $v_1^{(i)}, \dots, v_\ell^{(i)}$  such that  $\text{rt}^{(i+1)} = \mathcal{H}(v_1^{(i)}, \dots, v_\ell^{(i)})$  and  $\text{rt}^{(i)}$  is among the  $v^{(i)}$ -s (denoting  $v^{(0)}$  by  $\text{rt}^{(0)}$ ).

Our main challenge is to keep at bay the complexity of proving (and verifying) a hash of size  $\ell$  at each level. In order to go from this general construction to our final concrete one, there are three additional steps:

1. Going from generic hash-functions to towers of elliptic curves
2. Adding zero-knowledge (from hashes to commitments with “select-and-rerandomize”)
3. Going from  $d$  to 2 proofs by moving from towers to 2-cycles of elliptic curves.

We now elaborate on each. We stress that, for sake of clarity, we somewhat simplify our explanation and leave out several optimizations we carry out in our final construction. See Section 3 and Section 4 for the actual construction.

**Efficient proofs through “tower hashing”** In order to obtain an efficient proof of hashing, we would like to apply a hash that is SNARK-native in the sense defined above. Our target will be applying a simple Pedersen vector hashing that is native for Bulletproofs, which is a transparent proof system. We, however, quickly run into an issue: this does not work for more than one level because it is not *structure-preserving*. An intuition about what we mean by that is: hashing at a single level would be no problem but when, when applied at multiple levels, we would need to hash the output of an earlier hash function. This would not be “native” anymore for the proof system. The problem arises because a Pedersen hash maps *field elements* to *group elements* and because we have multiple levels: a parent  $x$  of a node—an hash image—will have to be hashed again to produce its own parent—thus being part of a preimage. Node  $x$  would need to be a point in the field and in the group at the same time.

In a general version of our solution, we solve this problem by using a different hash function at each level so that we can prove it natively inside Bulletproofs at each level. In order to do this we exploit a tower of curves, with a curve at each different level. We provide more details in Section 3.1. Through this solution we can simply produce a root  $\text{rt}^{(i+1)}$  of  $\ell$  children  $v_1^{(i)}, \dots, v_\ell^{(i)}$  by representing each child as a pair of coordinates  $(\mathbf{x}, \mathbf{y}) \in \mathbb{F}_p$  in the base field and producing a Pedersen hash with  $2\ell$  generators in  $\mathbb{E}_q(\mathbb{F}_p)$ . We can thus produce a group element  $H$  lying in  $\mathbb{E}_q(\mathbb{F}_p)$  and represent it (and its siblings) as pairs in  $\mathbb{F}_q \times \mathbb{F}_q$ . At the next level we can do the same using  $2\ell$  generators in  $\mathbb{E}_r(\mathbb{F}_q)$  for another elliptic curve of order  $r$ . And so in the same way till we get to the root. In our concrete solution we will reduce this tower to a 2-cycle and use only two elliptic curves.

**Adding Zero-Knowledge** So far we have been concerned with solutions that do not hide which element in the vector we are opening. We actually provided the internal nodes we encounter along the opening path; therefore, in order to make our solution private, we need to modify it so to provide a *masking* of each of the internal nodes along the opening path. That is, instead of sending the actual internal node  $\text{rt}^{(i)}$  (the hash of its children), we let the prover sample some fresh randomness  $[\rho']$  and send rerandomization (a commitment)  $\text{cm}^{(i)} = \text{rt}^{(i)} + [\rho'] \cdot H$  (where  $h$  is a group generator). What the verifier should be shown *in zero-knowledge* at each level is a slight variation on the relation we considered before. For level  $i$ , Given  $\text{cm}^{(i+1)}$  and  $\text{cm}^{(i)}$  (respectively, the rerandomized nodes along the path at the next and current level) the verifier should be guaranteed that the prover knows some  $v_1^{(i)}, \dots, v_\ell^{(i)}, [\rho], [\rho']$  and index  $j \in \{1, \dots, \ell\}$  such that (a)  $\text{cm}^{i+1} = \mathcal{H}(v_1^{(i)}, \dots, v_\ell^{(i)}, [\rho'])$  (the hash is again a Pedersen hash but now randomized through  $[\rho']$ ) and (b)  $\text{cm}^{(i)} = v_j^{(i)} + [\rho]H$ . What is important for efficiency is that relation (a)—a multiscalar exponentiation—can again be proved natively as before. Relation (b)—a single multiplication—needs to be expressed as an arithmetic circuit. This is still relatively inexpensive for us since because it’s performed once per level in a shallow tree.

**Optimizing with a 2-cycle and further technical points** We optimize our proof size further by applying an observation. We can move from a tower of  $d$  curves to only 2-cycle (two curves only). We can use this to reduce our communication complexity since instead of having  $d$  proofs—one per level each with a different curve—we can

<sup>7</sup> We anticipate that our final solution will have reduce to two proofs instead of  $d$ .

produce together two proofs—each referring to  $d/2$  of the levels. See also Section 4.3 for other optimizations and technical points.

**Vcash: transparent practical anonymous transactions.** We use Curve Trees as a main building block in Vcash. We defer the reader to [CHA21] for some high level ideas on the architecture (see in particular Section 1.3). The basic idea is to have the unspent coins in the systems stored as a set  $S$  of commitments in an accumulator. At the moment of transferring a coin, a user would prove in zero-knowledge that they own the coin (that is, they know the opening of some element in the set  $S$ ) and provide a rerandomization of that coin (which also needs to be appropriately proved in zero-knowledge). We describe our construction at the high-level in Section 5.

## 1.2 Related Work

When comparing to existing approaches to zero-knowledge for set membership we focus on succinctness and prover efficiency.

Some works with transparent setup do not achieve succinctness (that is, practically short proofs and a  $o(|S|)$  verification time). For example, Monero [AJ18]—or, generally, approaches based on ring signatures—have proofs linear in the set and where the verifier’s running time is linear in the size of the set  $|S|$ .

Other approaches to accumulator with zero-knowledge properties do not involve general-purpose SNARKs. This includes for example the multilinear pairing-based polynomial commitment in [BCF<sup>+</sup>21], the seminal KZG [KZG10] and the polynomial commitments in [BMM<sup>+</sup>21]. They, however, all require knowledge-based assumptions and a trusted setup. Similar observations hold for the recent work in Caulk[ZBK<sup>+</sup>22].

Other works apply asymptotically efficient polynomial commitments with a transparent setup, but their commitment and proof size are concretely large. This is the case of Hyrax [WTs<sup>+</sup>18], where for large set sizes commitments can be  $\gg 10KB$ , and Dory [Lee21] where commitments are 190 bytes (6-7 times larger than ours). Proofs of single opening are also large (18 KB) in Dory, although the scheme can amortize this cost with batching (expect for very large opening batches this amortized proof size is still significantly higher than ours). The Spartan proof system has overall opening sizes, proving and verification time that are competitive with respect to ours (for sets up to approximately  $2^{20}$  where Spartan starts to perform worse), but it has very high commitment sizes, e.g.  $\geq 20KB$  for sets of size  $2^{20}$  ( $625\times$  worse than ours)<sup>8</sup>. Other transparent polynomial commitments include those based on Reed-Solomon IOPs [BBHR18] or on Diophantine ARGuments of Knowledge (DARK) [BFS20]. As argued in [Lee21] (Section 1.1) they achieve worse concrete performances than the works above in practice.

Works that apply specialized proving techniques on accumulators in unknown-order groups: Veksel, [CHA21, CHA22, BCF<sup>+</sup>21, CFH<sup>+</sup>21]. These works obtain concretely small proofs/verifier with an efficient proving time, but require an RSA modulus (non-transparent) for their efficient instantiations<sup>9</sup>. While the work in [CFH<sup>+</sup>21] obtains concretely efficient proving time with a slightly larger proof size in ZCash but it require trusted setup to instantiate its proof system in addition to RSA modulus.

Using “friendlier” hash functions for Merkle trees mitigates the complexity of proving an opening. One such example is Poseidon [GKR<sup>+</sup>21]. The limitation of these solutions is that they rely on hash functions which are quite new and have not received the proper cryptanalytic scrutiny yet.

COMPARING OUR ACCUMULATOR TO TRANSPARENT ALGEBRAIC MERKLE TREES. The most interesting comparison to our (zero-knowledge) accumulator construction is a “transparent” version of that used in ZCash. Here, for to show membership in a set we apply a specific type of Merkle tree. In it, the collision-resistant hash function we use at each level has an extra property and in particular we require it to be “algebraic” (that is it can be expressed as a polynomial in a ring). A natural choice for this—and the one applied in ZCash—is Pedersen hash. In order to prove membership we apply a zkSNARK to the opening of the Merkle tree. For this approach to be efficient we need that the group in which we compute the hash is tied to the group in which the zkSNARK “functions”<sup>10</sup>. ZCash uses Groth16 on curve BLS12-381 [Gro16] as a zkSNARK and a specific curve for hashing, JubJub. Nonetheless, this approach could be made transparent by applying Bulletproofs on the Ristretto curve and choosing an appropriate elliptic curve for hashing (this includes for example Jabberwock in [CHA21]). In the remainder of this comparison we refer to this way of transparently instantiating Merkle trees with Pedersen hash as AlgMT<sub>BP</sub>.

<sup>8</sup> See [Lee21] for numbers referred in this section.

<sup>9</sup> In all these works we can replace the RSA group with a transparent class-group [BH01], substantially increasing their complexity. See, e.g., discussion in [DGS20]

<sup>10</sup> More specifically, this means that elliptic curve of the zkSNARK should be of order related to that of the definition field of the elliptic curve we use for Pedersen hashing.

We now compare the approach in our work to that in  $\text{AlgMT}_{\text{BP}}$ . Let us denote by  $N$  the set size. First we observe that asymptotically  $\text{AlgMT}_{\text{BP}}$  requires performing  $\log_2(N)$  hash computations with 2 elements each (one per level of tree, hence  $\log_2(N)$ ) inside the SNARK. Concretely, for a representative choice of  $N = 2^{32}$  to  $\approx 45000$  constraints. Our construction, on the other hand, requires  $\approx 5000$  constraints (almost an order of magnitude less).

COMPARING TO VERKLE TREES. At the *very* high-level, our approach resembles the currently explored “Verkle Tree” (VT) approach in Ethereum<sup>11</sup>. In both approaches an internal node represents a vector commitment to its children. The two approaches have a few substantial differences. First, that approach is currently not *structure-preserving* in the sense ours is. Each node is a commitment to a *hash* (e.g., SHA or Blake) of the children. This is required to solve a similar problem to that we approach with towers of curves. Currently Verkle Trees do no account for zero-knowledge (our focus in this work). If they did, we would also need to show in zero-knowledge that hashing the children has been performed correctly. For example, in the case of Blake this would require 20K additional constraints per level<sup>12</sup> (our solution, on the other hand, is in the ballpark of 5K constraint *in total*). Another difference is that the current implementation uses KZG polynomial commitments [KZG10] which requires a trusted setup.

CURVE TREES AND HALO2. Halo2<sup>13</sup> is a concrete transparent (zero-knowledge) proof system that uses recursion. It is concretely efficient and it obtains recursion by going back and forth in a cycle of two curves. Halo2 and curve trees have orthogonal goals: one is a full-fledged proof system, the other can be seen as a specialized data structure (and related constraint system) for zero-knowledge for set membership. We see, however, great potential in combining the techniques in these two systems and we are currently working in this direction.

## 2 Preliminaries

### 2.1 Notation

We denote by  $\mathbb{E}[\mathbb{F}]$  the elliptic curve  $\mathbb{E}$  defined over the field  $\mathbb{F}$ , whenever clear from context we might omit the field of definition and simply write  $\mathbb{E}$ . Whenever possible we explicitly mark scalar elements through square brackets and group elements with upper case letters (we will occasionally break this convention if not important). We denote by  $\mathfrak{M}(\mathcal{R})$  an upper bound (by construction) on the multiplicative complexity of verifying the relation  $\mathcal{R}$ . In practice, when estimating performance, the latter is the primary metric when proving satisfiability of arithmetic circuits.

When expressing an NP relation  $R(x, w)$  we make the private witness  $w$  explicit as such but we keep the public statement  $x$  implicit. For example, in the relation  $\mathcal{R}$  below

$$\mathcal{R} := \{z : y = \text{SHA}(g^z)\}$$

the only private witness is  $z$ , while  $g$  and  $y$  are considered publicly known inputs.

### 2.2 Commitments

We use the following syntax for commitments:

**Definition 1 (Commitments).** *A commitment scheme  $\mathcal{C}$  is a pair of algorithms (Setup, Comm) with syntax:*

- $\text{Setup}(1^\lambda) \rightarrow \text{ck} : \text{generates a commitment key ck};$
- $\text{Comm}(\text{ck}, m; r) \rightarrow c_m : \text{produces commitment } c_m \text{ to message } m \text{ with randomness } r.$

As it is standard, we call *message space* the set of  $m$ -s for which  $\text{Comm}$  is defined and commitment space its range,  $\text{Rng}(\text{Comm})$ . We require commitments to be *perfectly hiding*—the distribution of  $\text{Comm}(\text{ck}, m; r)$  is identical to the uniform distribution over the commitment space—and *computationally binding*—no efficient adversary can produce two pairs  $(m, r), (m', r')$  such that  $m \neq m'$  and  $\text{Comm}(\text{ck}, m; r) = \text{Comm}(\text{ck}, m'; r')$ .

*Remark 1 (Rerandomizable Commitments).* We will use rerandomizable commitments, i.e., endowed with an algorithm  $\text{Rerand}(\text{ck}, \text{Comm}(\text{ck}, m; r)) \rightarrow (c', r')$  such that  $c' = \text{Comm}(\text{ck}, m; r + r')$ . Notice that homomorphic commitments (and thus Pedersen commitments) satisfy this property.

<sup>11</sup> <https://dankradfeist.de/ethereum/2021/06/18/verkle-trie-for-eth1.html>

<sup>12</sup> <https://github.com/zcash/zcash/issues/2258>

<sup>13</sup> <https://electriccoin.co/blog/explaining-halo-2/>

### 2.3 Accumulators

**Definition 2 (Accumulator scheme).** An accumulator scheme  $\text{Acc}$  over universe  $\mathcal{U}_\lambda(\text{Acc})$  (where  $\lambda$  is a security parameter) consists of PPT algorithms  $\text{Acc} = (\text{Setup}, \text{Accum}, \text{PrvMem}, \text{VfyMem})$  with the following syntax:

$\text{Setup}(1^\lambda) \rightarrow (\text{pp})$  generates public parameters  $\text{pp}$ .

$\text{Accum}(\text{pp}, S) \rightarrow A$  deterministically computes accumulator  $A$  for set  $S \subseteq \mathcal{U}_\lambda(\text{Acc})$ .

$\text{PrvMem}(\text{pp}, S, x) \rightarrow W$  computes witness  $W$  that proves  $x$  is in accumulated set  $S$ .

$\text{VfyMem}(\text{pp}, A, x, W) \rightarrow b \in \{0, 1\}$  verifies through witness whether  $x$  is in the set accumulated in  $A$ . We do not require parameter  $x$  to be in  $\mathcal{U}_\lambda(\text{Acc})$  from the syntax.

An accumulator scheme should satisfy correctness—the accumulator works as expected—and soundness—no efficient adversary can choose a set  $S$  and then find a witness that checks on  $\text{Acc.Accum}(\text{pp}, S)$  and  $x \notin S$ <sup>14</sup>.

### 2.4 NIZKs

Non-Interactive Zero-Knowledge schemes (or NIZKs) require a reference string which can be either uniformly sampled (a *urs*), or structured (a *srs*). In the latter case it needs to be sampled by a trusted party. In this work we use and assume *transparent* NIZKs, i.e. whose algorithms use a reference string *urs* sampled uniformly.

**Definition 3.** A NIZK for a relation family  $\mathfrak{R} = \{\mathfrak{R}_\lambda\}_{\lambda \in \mathbb{N}}$  is a tuple of algorithms  $\text{ZK} = (\text{Prove}, \text{VerProof})$  with the following syntax:

- $\text{ZK.Prove}(\text{urs}, R, x, w) \rightarrow \pi$  takes as input a string *urs*, a relation description  $R$ , a statement  $x$  and a witness  $w$  such that  $R(x, w)$ ; it returns a proof  $\pi$ .
- $\text{ZK.VerProof}(\text{urs}, R, x, \pi) \rightarrow b \in \{0, 1\}$  takes as input a string *urs*, a relation description  $R$ , a statement  $x$  and a proof  $\pi$ ; it accepts or rejects the proof.

We require a NIZK to be complete, that is, for any  $\lambda \in \mathbb{N}$ ,  $R \in \mathfrak{R}$  and  $(x, w) \in R$  it holds with overwhelming probability that  $\text{VerProof}(\text{urs}, R, x, \pi)$  where  $\text{urs} \leftarrow_{\$} \{0, 1\}^{\text{poly}(\lambda)}$  and proof  $\pi \leftarrow \text{Prove}(\text{urs}, R, x, w)$ .

We also require knowledge-soundness and zero-knowledge to hold. Informally, the former states we can efficiently “extract” a valid witness from a proof that passes verification; the latter states that the proof leaks nothing about the witness (this is modeled through a simulator that can output a valid proof for an input in the language without knowing the witness). We use variants of these notions with certain composability properties, e.g. requiring auxiliary inputs and relation generators. For a full formal treatment of these, we refer the reader to Sections 2.2 and 2.5 in [BCFK19].

Whenever the relation family is obviously defined, we talk about a “NIZK for a relation  $R$ ”.

*Remark 2 (Relations and Public Inputs).* In the algorithms above we have both a relation  $R$  and a public input  $x$  as inputs. The reason is that in a soundness experiment,  $R$  may be constrained to be from a certain distribution on  $\mathfrak{R}$  whereas  $x$  can be chosen arbitrarily by the adversary. See for example Section 2.2 in [BCFK19]. In our constructions we often assume prover and verifier to implicitly take as input the relation description<sup>15</sup>.

In the proof of security of our construction we require an additional property for one of our NIZKs, *simulation-extractability*. Namely, extractability should hold even with respect to an adversary that has access to simulated proofs. We refer the reader to [Gro06] for formal definitions.

**Modular NIZKs through Commit-and-Prove.** We use the framework for black-box modular composition of commit-and-prove NIZKs (or CP-NIZKs) in [CFQ19] and [BCFK19]. Informally a CP-NIZK is a NIZK that can efficiently prove properties of committed inputs through some commitment scheme  $\text{C}$ . Let  $x$  be a public input and  $c$  a commitment. Such a scheme can for example prove knowledge of  $(u, \omega, r)$  such that  $c = \text{Comm}(u; r)$  and that relation  $R_{\text{inner}}(x; u, \omega)$  holds. We can think of  $\omega$  as a non-committed part of the witness. Besides the proof, the verifier’s inputs are  $x$  and  $c$ .

In our construction we will make use of the following folklore composition to obtain efficient NIZKs from CP-NIZKs. Fixed a commitment scheme and given two CP-NIZKs  $\text{CP}, \text{CP}'$  respectively for two “inner” relations  $R$  and  $R'$ , we can prove their conjunction (for a shared witness  $u$ )  $R^*(x, x', u, \omega, \omega') = R(x, u, \omega) \wedge R'(x', u, \omega')$  like this: the prover commits to  $u$  as  $c \leftarrow \text{Comm}(u, r)$ ; generates proofs  $\pi$  and  $\pi'$  from the respective schemes; it outputs combined proof  $\pi^* := (c, \pi, \pi')$ . The verifier checks each proof over respective inputs  $(x, c)$  and  $(x', c')$ .

The following theorem (informally stated) is a direct consequence of Theorem 3.1 in [CFQ19].

<sup>14</sup> These definitions are standard and we refer the reader to [BBF19] for a formal treatment.

<sup>15</sup> This parameter is usually short.

**Theorem 1 (Black-Box Composition of CP-NIZKs).** *The construction above is a secure NIZK for the conjunction relation  $R^*$ .*

We can see Bulletproof [BBB<sup>+</sup>18] as a CP-NIZK since it works efficiently over an implicit commitment representation (see discussion in [CFQ19]). We use this fact in our instantiations.

### 3 Curve Trees as Accumulators

#### 3.1 Towers of Elliptic Curves

We call a sequence of elliptic curves  $\mathbb{E}_1(\mathbb{F}_{p_1}), \dots, \mathbb{E}_n(\mathbb{F}_{p_n})$  a tower, iff. for  $i \in [1, n-1] : p_i = |\mathbb{E}_{i+1}(\mathbb{F}_{p_{i+1}})|$ ; in other words the field of  $\mathbb{E}_i(\mathbb{F}_{p_i})$  is the scalar field of  $\mathbb{E}_{i+1}(\mathbb{F}_{p_{i+1}})$ . We will generally let the field of definition be implicit to simplify notation. Towers of curves have previously been used to optimize the proving of cryptographic operations in zkSNARKs [KZM<sup>+</sup>15] [HBHW21], which will also be the application in this paper. Additionally the same techniques has been applied for recursive proofs systems [BCTV14] [BGH19]. We will not require that any of the curves are pairing friendly.

**Cycles of Elliptic Curves.** Of particular interest are  $m$ -cycles of elliptic curves: infinitely long towers where  $\mathbb{E}_i(\mathbb{F}_{p_i}) = \mathbb{E}_{i+m}(\mathbb{F}_{p_{i+m}})$  for all  $i$ . Most commonly  $m = 2$ , which will be the primary case of interest in this paper as well.

#### 3.2 The Accumulator: $(\ell, \mathbb{E}_1, \dots, \mathbb{E}_d)$ -Curve Tree

Our accumulation scheme is a ‘‘Curve Tree’’: a Curve Tree can be seen as an ‘algebraically compatible’ Merkle tree which uses a Pedersen commitments over  $\mathbb{E}_1, \dots, \mathbb{E}_d$  as the compression function for each level respectively. Notice that below we use a ‘‘randomness’’ scalar  $[r]$ . This will be useful for several concrete aspects of our construction described in the next section.

**Definition 4 (Curve Trees).** *A Curve Tree, parameterized by a branching factor  $\ell$  and a tower of Elliptic curves  $\mathbb{E}_1, \dots, \mathbb{E}_d$  has the following recursive structure:*

**Leaf:**  $(\ell, \mathbb{E}_1)$ -CurveTree: *A  $(\ell, \mathbb{E}_1)$ -Curve Tree (leaf node) is a 3-tuple  $(C, 0, C)$  where  $C \in \mathbb{E}_1$ : a tree containing  $C$  with root  $C$ .*

**Parent:**  $(\ell, \mathbb{E}_1, \dots, \mathbb{E}_d)$ -CurveTree: *A  $(\ell, \mathbb{E}_1, \dots, \mathbb{E}_d)$ -Curve Tree is a 3-tuple  $(C, r, (T_1, \dots, T_\ell))$ , where  $T_1, \dots, T_\ell$  are  $(\ell, \mathbb{E}_1, \dots, \mathbb{E}_{d-1})$ -Curve Trees, i.e.*

$$T_1 = (\hat{C}_1 = (\mathfrak{x}_1, \mathfrak{y}_1), \hat{r}_1, \hat{T}_1), \dots, T_\ell = (\hat{C}_\ell = (\mathfrak{x}_\ell, \mathfrak{y}_\ell), \hat{r}_\ell, \hat{T}_\ell)$$

*The root  $C$  of the  $(\ell, \mathbb{E}_1, \dots, \mathbb{E}_d)$ -Curve Tree is a Pedersen commitment to the coordinates of the children:*

$$C = \langle \mathfrak{x}, \vec{G}^{(d,X)} \rangle + \langle \mathfrak{y}, \vec{G}^{(d,Y)} \rangle + [r] \cdot H^{(d)}$$

*For constants  $\vec{G}^{(d,X)}, \vec{G}^{(d,Y)} \in \mathbb{E}_d^\ell$  and  $H^{(d)} \in \mathbb{E}_d$ .*

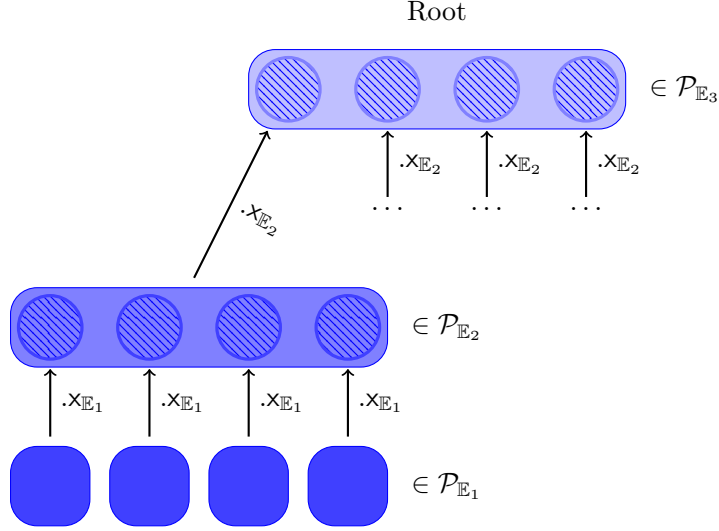
#### 3.3 Supported Curves.

Since we only require the hardness of discrete log, the techniques described in this paper are broadly applicable to any tower of curves and in particular existing cycles of elliptic curves, e.g. the ‘‘Pasta Cycle’’ (Pallas / Vesta curves) [Hop20], the ‘‘Tweedle Cycle’’ (Tweedledum / Tweedledee curves) [Hop19] or the Secp256k1 / Secq256k1 curves cycle [Poe18]. Even though our techniques do not use bilinear pairing, security of our zero-knowledge accumulator holds even in the presence of an efficiently computable bilinear map (type I, type II or type III) on one/both of the curves – enabling interoperability of our techniques with proof systems over such cycles.

### 4 Zero-Knowledge Set Membership in Curve Trees

In this section we describe how to prove set membership in zero knowledge for our curve trees accumulators. We use a slightly more specific version of the relation for zero knowledge for set membership in, e.g., [BCF<sup>+</sup>21, CFH<sup>+</sup>21]. Our variant (dubbed ‘‘Select-and-Rerandomize’’) obtains stronger unlinkability properties and is described below. We stress that our construction can be straightforwardly adapted to fit the setting in [BCF<sup>+</sup>21, CFH<sup>+</sup>21] where the commitment to the element is already given and there is no need for rerandomization.

Our final scheme achieves  $O(\log n)$  communication and  $O(\sqrt[d]{n})$  computation where  $d$  is the depth of the tree and  $n$  is the size of the accumulated set.



**Fig. 1.** Illustration of a  $(4, \mathbb{E}_1, \mathbb{E}_2, \mathbb{E}_3)$ -Curve Tree. Hatch pattern circles indicates that the point is represented as a field element, rounded rectangles represents Pedersen commitments to the field elements (circles) inside. Darker shades indicates lower levels in the tree.

#### 4.1 “Select-and-Rerandomize” Accumulators

Here we define an auxiliary interface for a primitive we call a “select-and-rerandomize accumulator”. Given an accumulator of committed values, I can provide you with a handle (a commitment) to a value in the accumulated set proving to you it is in the set but without revealing which one it is. We achieve this through rerandomization of commitments and zero-knowledge proofs over the set membership relation and commitment rerandomization. This primitive is natural in several settings, such as in anonymous payment systems, including the one in this work and in [CHA21]. Similar attempts at having hiding for accumulator witnesses (with different techniques) also appeared in [CFH<sup>+</sup>21].

Below we assume an accumulator scheme and an accumulated set  $S$  whose elements are (rerandomizable) commitments. We also denote by  $\text{pp}$  the concatenation of the accumulator parameters and the commitment key.

$\text{SelRerand}.\mathcal{P}(\text{pp}, S, c) \rightarrow (c', \pi, r')$  returns a rerandomized commitment (of  $c \in S$ ), a proof of membership and the (auxiliary) randomness used for rerandomization.

$\text{SelRerand}.\mathcal{V}(\text{pp}, A, c', \pi) \rightarrow 0/1$  verifies that  $c'$  is a rerandomization of an element in the set.

**Correctness:** For any set  $S$ ,  $c \in S$  (with  $c$  an honestly generated commitment), honestly generated parameters  $\text{pp} = (\text{pp}_{\text{acc}}, \text{ck})$  the following holds

$$\text{SelRerand}.\mathcal{V}(\text{pp}, \text{Accum}(\text{pp}_{\text{acc}}, S), c', \pi) = 1 \wedge c' = \text{Rerand}(\text{ck}, c; r')$$

where  $\text{SelRerand}.\mathcal{P}(\text{pp}, S, c) \rightarrow (c', \pi, r')$ .

**Security (informal):** we require the proof  $\pi$  to be an extractable NIZK (i.e., we require knowledge soundness and zero-knowledge) for the relation below:

$$\mathcal{R}^{\text{(SelectRerand)}} := \left\{ (W, c, r) : \begin{array}{l} c' = \text{Rerand}(\text{ck}, c; r') \\ \wedge \text{Acc}.\text{VfyMem}(\text{pp}_{\text{acc}}, A, c, W) \end{array} \right\}$$

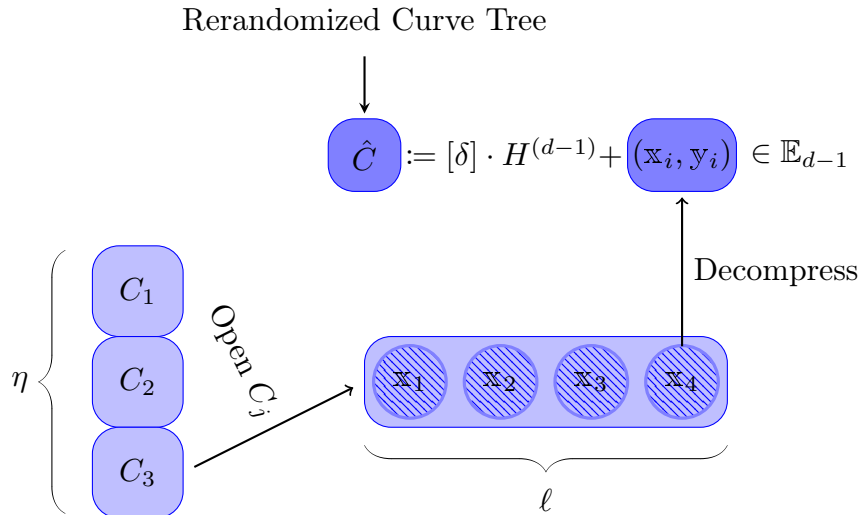
whenever the parameters and the accumulator have been generated honestly.

#### 4.2 Constructing Select-and-Rerandomize in Curve Trees

The main intuition is as follows. Observe that the leafs of a  $(\ell, \mathbb{E}_1, \dots, \mathbb{E}_d)$ -Curve Tree are curve points on  $\mathbb{E}_1$ . These leafs are going to be Pedersen commitments to secret vectors. We exploit the recursive algebraic structure of our tree to enable efficient zero-knowledge proofs of knowledge of these vectors. A crucial feature of Curve Trees is that the roots (and that of every sub-tree) are rerandomizable commitments to Curve Trees: by rerandomizing the root



$C$  as  $C' \leftarrow C + [\delta] \cdot H^{(d)}$  we obtain a perfectly hiding commitment to the same set of children as the original tree, we exploit this observation to traverse the tree level-by-level using a simple zero-knowledge proof described in the next section. Another property of curve trees is that the height can dynamically increase as the number of elements in the tree grows: by using a cycle of curves the tree can grow “upward” while the leaves remain  $\mathbb{E}_1$  points.



**Fig. 2.** Illustration of the “select and rerandomize relation”. In the example illustrated above  $i = 4$  and  $j = 3$ . Note that  $\hat{C}$  is one layer deeper in the tree.

**4.2.1 Single-Level Select-and-Rerandomize** The central component in our construction is a simple construction for a select-and-rerandomize-like relation for a *single* level in a curve tree (we later recursively invoke this to obtain a full select-and-rerandomize in Section 4.2.2). Its underlying relation takes as input a rerandomized commitment  $\hat{C}$ , an alleged parent  $C$  and hidden indexes (the root of a  $(\ell, \mathbb{E}_1, \dots, \mathbb{E}_d)$ -Curve Tree), secret index  $i$  whose semantics is “ $\hat{C}$  is the  $i$ -th child of  $C$ ”. Practically, this is accomplished by opening the commitment  $C$  to  $\vec{x}, \vec{y}$  using a Pedersen commit-and-prove over  $\mathbb{E}_d$ , then rerandomizes the  $i$ ’th point  $\hat{C} = (\mathbf{x}_i, \mathbf{y}_i) + [\delta] \cdot H^{(d-1)} \in \mathbb{E}_{d-1}$ . Slightly more formally we require a zero-knowledge argument of knowledge for the following relation:

$$\mathcal{R}^{(1\text{-Lvl-CrvT})} := \left\{ (i, r, \delta, \vec{x}, \vec{y}) : \begin{array}{l} C = \langle \vec{x}, X^{\vec{t}} \rangle + \langle \vec{y}, Y^{\vec{t}} \rangle + [r] \cdot H^{(d)} \\ \wedge \hat{C} = (\mathbf{x}_i, \mathbf{y}_i) + [\delta] \cdot H^{(d-1)} \end{array} \right\}$$

As noted the  $C = \langle \vec{x}, X^{\vec{t}} \rangle + \langle \vec{y}, Y^{\vec{t}} \rangle + [r] \cdot H^{(d)}$  constraint can be very efficiently enforced using a commit-and-prove for Pedersen commitments (e.g. Bulletproofs or Compressed  $\Sigma$ -Protocol) over  $\mathbb{E}_d$ . While  $\hat{C} = (\mathbf{x}_i, \mathbf{y}_i) + [\delta] \cdot H^{(d-1)}$  requires a single fixed-based exponentiation “inside the circuit”. We describe an optimized arithmetic circuit for the relation above in Appendix B.

In the appendix we describe a generalization of this relation to the “forest” case where we do not have only one root but several roots (this can be used for a flatter construction described in Remark 3).

**4.2.2 Recursive  $\sqrt[d]{n}$  Membership Proof.** We can observe that  $\mathcal{R}^{(1\text{-Lvl-CrvT})}$  already yields a “one-level” select-and-rerandomize. We can apply it one level at the time and obtain a full construction. Note that in the single-level case, at a given level we can provide a (hiding) Pedersen commitment of one of the children. The latter in turn represents the root of a (sub-)Curve Tree. We can thus extend the technique presented to obtain membership proofs with  $O(\sqrt[d]{n})$  prover/verifier complexity: by letting  $\ell = \sqrt[d]{n}$  and using  $d - 1$  individual proofs and  $\eta = 1$  (except possibly for the first proof which sets  $\eta = \sqrt[d]{n}$ ). The proofs then descends down the Curve Tree one level at a time.

**Theorem 2 (informal).** *The construction in Section 4.2.2 is select-and-rerandomize over curve trees as accumulators with  $O(\sqrt[d]{n})$  prover/verifier complexity. This construction is transparent if we instantiate the underlying NIZK with Bulletproofs.*

SelRerand. $\mathcal{P}(\text{pp}, S, c^*)$	SelRerand. $\mathcal{V}(\text{pp}, \text{rt}, c^*, \pi^*)$
Reconstruct tree from $S$ ; let $\text{rt}$ be its root Let $c^{(0)}, \dots, c^{(d)}$ be the path elements to $c^*$ in the tree (with $c^{(0)}$ corresponding to $\text{rt}$ ) Let $\hat{c}^{(0)} := \text{rt}$ and $r^{(0)} := 0$ <b>for</b> $j = 1, \dots, d$ <b>do</b> $(\hat{c}^{(j)}, r^{(j)}) \leftarrow \text{Rerand}(\text{ck}, c^{(j)})$ $\pi_j \leftarrow \text{ZK.Prove}(\text{pp}, \mathcal{R}^{(1\text{-Lvl-CrvT})}, \hat{c}^{(j-1)}, \hat{c}^{(j)}, r^{(j)}, r^{(j-1)})$ <b>endfor</b> Return $\pi^* := (\hat{c}^{(1)}, \dots, \hat{c}^{(d)}, \pi^{(1)}, \dots, \pi^{(d)})$	Parse $\pi^*$ as $(\hat{c}^{(1)}, \dots, \hat{c}^{(d)}, \pi^{(1)}, \dots, \pi^{(d)})$ Let $\hat{c}^{(0)} := \text{rt}$ <b>for</b> $j = 1, \dots, d$ <b>do</b> $b_j \leftarrow \text{ZK.VerProof}(\text{pp}, \mathcal{R}^{(1\text{-Lvl-CrvT})}, \hat{c}^{(j-1)}, \hat{c}^{(j)})$ <b>endfor</b> Accept iff $\bigwedge_{j=1, \dots, \ell} b_j = 1$

**Fig. 3.** Select-and-rerandomize for a  $(\ell, \mathbb{E}_1, \dots, \mathbb{E}_d)$ -CurveTree

The theorem above can be proved straightforwardly by invoking Theorem 1 at each level in the tree. Zero-knowledge follows straightforwardly by the rerandomization of the commitments and the zero-knowledge of the underlying NIZK.

### 4.3 Optimizations

In this section, we briefly cover some straightforward optimizations.

**Merging Proofs.** The approach above which necessitates a total of  $d - 1$  individual proofs  $\pi_1, \dots, \pi_{d-1}$ . However, when using a (2) cycle of curves, all the even/odd proofs are over the same curve/field and can therefore be combined in to a single larger statement; only requiring 2 proofs in total for any  $d$ .

**Point Compression / Permissible Points.** In the setting where the accumulator is guaranteed to have been computed honestly (e.g. in our confidential transactions application), we can reduce the number of exponentiations during committing and the size of the witness by only committing to the  $\mathbf{x}$ -coordinate of the children: this remains binding by ensuring that only one of  $(\mathbf{x}, \mathbf{y})$  and  $(\mathbf{x}, -\mathbf{y})$  is “allowed”. One common choice is to take the numerically smallest between  $\mathbf{y}$  and  $-\mathbf{y}$ , or discriminate based upon the parity (even/odd) over  $\mathbb{Z}$ , however neither of these constraints can be efficiently expressed as an arithmetic circuit; instead we use a universal hash function. Let  $S(v) = 1$  iff.  $v \in \mathbb{F}$  is a quadratic residue (i.e. there exists  $w \in \mathbb{F}$  st.  $w^2 = v$ ) and  $S(v) = 0$  otherwise. Now consider the following family of 2-universal hash functions from any field to  $\{0, 1\}$ :

$$\begin{aligned} \mathcal{U}_{\alpha, \beta}(v) &: \mathbb{F} \rightarrow \{0, 1\} \\ \mathcal{U}_{\alpha, \beta}(v) &\mapsto S(\alpha \cdot v + \beta) \end{aligned}$$

Observe that the constraint  $\mathcal{U}_{\alpha, \beta}(v) = 1$  can be enforced using a circuit with multiplicative complexity 1, showing  $\{(w) : w^2 = (\alpha \cdot v + \beta)\}$ . We exploit this to efficiently define a set of “permissible points” on  $\mathbb{E}$ :

$$\mathcal{P}_{\mathbb{E}} = \{(\mathbf{x}, \mathbf{y}) \mid (\mathbf{x}, \mathbf{y}) \in \mathbb{E}(\mathbb{F}_p) \wedge \mathcal{U}_{\alpha, \beta}(\mathbf{y}) = 1 \wedge \mathcal{U}_{\alpha, \beta}(-\mathbf{y}) = 0\}$$

Note that  $1/4$  of the points on  $\mathbb{E}$  are permissible and any  $(\mathbf{x}, \mathbf{y}) \in \mathcal{P}_{\mathbb{E}}$  is uniquely defined by its  $\mathbf{x}$ -coordinate – this is the case for any finite field of characteristic  $\notin \{2, 3\}$ . Any Pedersen commitment  $C$  can be “made permissible” by simply adding  $H$  (“incrementing the randomness”) until the point is permissible:

$$\begin{aligned} &\text{MakePermissible}(C) : \mathbb{E} \rightarrow \mathcal{P}_{\mathbb{E}} \\ &\text{1: } \mathbf{while } C \notin \mathcal{P}_{\mathbb{E}} : C \leftarrow C + H \\ &\text{2: } \mathbf{return } C \end{aligned}$$

In expectation, this requires 4 curve additions and 8 square roots. By enforcing that only permissible points are added to the accumulator<sup>16</sup> so that the “decompression” is unique, we can reduce the complexity of  $\mathcal{R}^{\text{(SelectRerand)}}$  slightly, instead showing:

$$\mathcal{R}^{(1\text{-Lvl-CrvT})} := \left\{ \begin{array}{l} C = \langle \vec{x}, X^{\vec{t}} \rangle + [r] \cdot H^{(d)} \\ (i, r, \delta, \vec{x}, \mathbb{y}) : \quad \wedge (\mathbb{x}_i, \mathbb{y}) \in \mathcal{P}_{\mathbb{E}^{(t)}} \\ \wedge \hat{C} = (\mathbb{x}_i, \mathbb{y}) + [\delta] \cdot H^{(d-1)} \end{array} \right\}$$

Note the  $(\mathbb{x}_i, \mathbb{y}) \in \mathcal{P}_{\mathbb{E}^{(t)}}$  constraint only requires a check that  $(\mathbb{x}_i, \mathbb{y}) \in \mathbb{E}^{(t)}$  in addition to  $\mathcal{U}_{\alpha, \beta}(\mathbb{y}) = 1$ . We include the full circuit description in Appendix B.

## 5 VCash: Transparent and Efficient Anonymous Payment System

In this section we informally describe our anonymous payment system, which we dub VCash. The techniques and model here follow mostly prior work; we defer extended and formal details to the final version of this paper.

Like ZCash our construction supports the largest possible anonymity set at every transaction. On the other hand, our scheme has an additional leakage: a party  $S$  sending a transaction tx to a party  $R$  can learn when  $R$  will spend the coins received in tx (but not to whom). Only sender  $S$  can infer this.

### 5.1 Model

The flow of our protocol roughly follows known blueprints. We defer the reader to Section 3 and 6 and Appendix D in [CHA21, CHA22] for a formal description of a closely related model. One important difference with the formal description in *Veksel*, however, is that we do not assume parties to hold accounts. Instead, in VCash (as in Zcash) a transaction pours pairs of input coins into pairs of output coins of equivalent value.

**Intuition about the model:** at any given moment in time, each party holds a certain number of coins<sup>17</sup> Each user is also holding a state roughly containing all the transfers occurred so far. Through the state, any user can verify the validity of each transfer and used to verify the validity of each.

**Components of the model:**

**Setup** The setup algorithm produces the initial parameters of the system. We emphasize that it does not require to be run by a trusted setup.

**Pouring** A sender  $S$  can “pour” the value of two *input coins* into two new *output coins* nullifying the input ones. The recipients of the two new coins can be distinct. It is possible for  $S$  itself to be one or both of the recipients. We require that the total value of input and output coins is the same.

**Verifying and Processing** We provide a verifying algorithm by which parties can check a transfer is valid (roughly that the sender could afford it) and processing algorithm by which parties can update their state after a transfer. Although we describe them separately for clarity they are run together in our formal construction in ??.

### 5.2 Construction

Again we defer the reader to the technical overview and Section 3 in [CHA21, CHA22] for further background. Here we provide high-level details.

A transaction consist of the creation of output coins from input coins. A coin roughly consists of a commitment to its amount and other information that ensures it will be used only once and by its intended recipient.

For a transaction to be valid it must be the case that:

1. Output coins are in an appropriate non-negative range (we want to *give* money and not take it in a transaction);
2. the total value of input and output coins is the same;

<sup>16</sup> In case of our anonymous cryptocurrency application, this is enforced by the network of block validators: as a condition for a transaction being valid.

<sup>17</sup> “Holding” a coin requires knowing a certain secret key associated with the user. In this section we ignore the aspect of registering with a new key to the system, but we stress it is straightforward to add.

3. input coins “exist” and are valid themselves.

We use zero-knowledge proofs to ensure the above. The first two properties can be ensured by range proofs and homomorphic properties of Pedersen commitments respectively. The last property is where we can use our select-and-rerandomize constructions from the previous sections. All coins are stored in an accumulator. When I want to spend an input coin I can select-and-rerandomize it obtaining a rerandomized version of that coin. Now I can include that in my transaction to prove this is the rerandomization of something existing in the accumulator.

Other aspects of the system (e.g., nullifiers) can be implemented straightforwardly (e.g. through public key rerandomization) and have small impact on the concrete efficiency of our construction; we *do* include these costs in our evaluation.

### 5.3 Concrete Efficiency

In VCash, the constraint system used for the zero-knowledge proof of a “spend” transaction is 20x smaller than that in ZCash Sapling (the currently deployed version of ZCash).

Concretely for two inputs/two outputs and anonymity sets of  $2^{32}$  (like in Zcash) our confidential transactions (Vcash) require participants to compute/verify two Bulletproofs proofs of  $\approx 5000$  constraints each. Verifying each of the proofs in parallel (2 cores) in batches of at least 100 transactions (e.g. when verifying the validity of all transactions in a block) yields a very practical per-transaction verification time of  $< 10$  ms. Transaction sizes are  $< 3$  KB. Our timings are competitive with those of the approach in ZCash Sapling and trade slightly larger proofs (proofs in ZCash are 0.2KB) for a completely transparent setup and simpler curve requirements.

## References

- ACD<sup>+</sup>16. Masayuki Abe, Melissa Chase, Bernardo David, Markulf Kohlweiss, Ryo Nishimaki, and Miyako Ohkubo. Constant-size structure-preserving signatures: Generic constructions and simple assumptions. *Journal of Cryptology*, 29(4):833–878, October 2016.
- AJ18. Kurt M. Alonso and Jordi Herrera Joacomartí. Monero - privacy in the blockchain. Cryptology ePrint Archive, Report 2018/535, 2018. <https://eprint.iacr.org/2018/535>.
- BBB<sup>+</sup>17. Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. Bulletproofs: Short proofs for confidential transactions and more. Cryptology ePrint Archive, Report 2017/1066, 2017. <https://eprint.iacr.org/2017/1066>.
- BBB<sup>+</sup>18. Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Greg Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *2018 IEEE Symposium on Security and Privacy*, pages 315–334. IEEE Computer Society Press, May 2018.
- BBF19. Dan Boneh, Benedikt Bünz, and Ben Fisch. Batching techniques for accumulators with applications to IOPs and stateless blockchains. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part I*, volume 11692 of *LNCS*, pages 561–586. Springer, Heidelberg, August 2019.
- BBHR18. Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Fast reed-solomon interactive oracle proofs of proximity. In Ioannis Chatzigiannakis, Christos Kaklamanis, Dániel Marx, and Donald Sannella, editors, *ICALP 2018*, volume 107 of *LIPICs*, pages 14:1–14:17. Schloss Dagstuhl, July 2018.
- BCF<sup>+</sup>21. Daniel Benarroch, Matteo Campanelli, Dario Fiore, Kobi Gurkan, and Dimitris Kolonelos. Zero-knowledge proofs for set membership: Efficient, succinct, modular. In Nikita Borisov and Claudia Diaz, editors, *Financial Cryptography and Data Security*, pages 393–414, Berlin, Heidelberg, 2021. Springer Berlin Heidelberg.
- BCFK19. Daniel Benarroch, Matteo Campanelli, Dario Fiore, and Dimitris Kolonelos. Zero-knowledge proofs for set membership: Efficient, succinct, modular. Cryptology ePrint Archive, Report 2019/1255, 2019. <https://eprint.iacr.org/2019/1255>.
- BCTV14. Eli Ben-Sasson, Alessandro Chiesa, Eran Tromer, and Madars Virza. Scalable zero knowledge via cycles of elliptic curves. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 276–294. Springer, Heidelberg, August 2014.
- BFS20. Benedikt Bünz, Ben Fisch, and Alan Szepieniec. Transparent SNARKs from DARK compilers. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 677–706. Springer, Heidelberg, May 2020.
- BGH19. Sean Bowe, Jack Grigg, and Daira Hopwood. Halo: Recursive proof composition without a trusted setup. Cryptology ePrint Archive, Report 2019/1021, 2019. <https://eprint.iacr.org/2019/1021>.
- BH01. Johannes Buchmann and Safuat Hamdy. A survey on iq cryptography. In *Public-Key Cryptography and Computational Number Theory*, pages 1–15, 2001.
- BMM<sup>+</sup>21. Benedikt Bünz, Mary Maller, Pratyush Mishra, Nirvan Tyagi, and Psi Vesely. Proofs for inner pairing products and applications. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 65–97. Springer, 2021.
- CFH<sup>+</sup>21. Matteo Campanelli, Dario Fiore, Semin Han, Jihye Kim, Dimitris Kolonelos, and Hyunok Oh. Succinct zero-knowledge batch proofs for set accumulators. Cryptology ePrint Archive, Report 2021/1672, 2021. <https://ia.cr/2021/1672>.
- CFQ19. Matteo Campanelli, Dario Fiore, and Anaïs Querol. LegoSNARK: Modular design and composition of succinct zero-knowledge proofs. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 2075–2092. ACM Press, November 2019.
- CHA21. Matteo Campanelli and Mathias Hall-Andersen. Veksel: Simple, efficient, anonymous payments with large anonymity sets from well-studied assumptions. Cryptology ePrint Archive, Report 2021/327, 2021. <https://ia.cr/2021/327>.
- CHA22. Matteo Campanelli and Mathias Hall-Andersen. Veksel: Simple, efficient, anonymous payments with large anonymity sets from well-studied assumptions. In *Proceedings of the 2022 ACM on Asia Conference on Computer and Communications Security*, pages 652–666, 2022.
- CHI<sup>+</sup>20. Megan Chen, Carmit Hazay, Yuval Ishai, Yuriy Kashnikov, Daniele Micciancio, Tarik Riviere, abhi shelat, Muthu Venkatasubramaniam, and Ruihan Wang. Diogenes: Lightweight scalable RSA modulus generation with a dishonest majority. Cryptology ePrint Archive, Report 2020/374, 2020. <https://eprint.iacr.org/2020/374>.
- DGS20. Samuel Dobson, Steven D. Galbraith, and Benjamin Smith. Trustless groups of unknown order with hyperelliptic curves. Cryptology ePrint Archive, Report 2020/196, 2020. <https://eprint.iacr.org/2020/196>.
- GKR<sup>+</sup>21. Lorenzo Grassi, Dmitry Khovratovich, Christian Rechberger, Arnab Roy, and Markus Schofnegger. Poseidon: A new hash function for zero-knowledge proof systems. pages 519–535. USENIX Association, 2021.
- Gro06. Jens Groth. Simulation-sound nizk proofs for a practical language and constant size group signatures. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 444–459. Springer, 2006.
- Gro16. Jens Groth. On the size of pairing-based non-interactive arguments. In Marc Fischlin and Jean-Sébastien Coron, editors, *EUROCRYPT 2016, Part II*, volume 9666 of *LNCS*, pages 305–326. Springer, Heidelberg, May 2016.

- GW20. Ariel Gabizon and Zachary J. Williamson. plookup: A simplified polynomial protocol for lookup tables. Cryptology ePrint Archive, Report 2020/315, 2020. <https://eprint.iacr.org/2020/315>.
- GWC19. Ariel Gabizon, Zachary J. Williamson, and Oana Ciobotaru. PLONK: Permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge. Cryptology ePrint Archive, Report 2019/953, 2019. <https://eprint.iacr.org/2019/953>.
- HBHW21. Daira Hopwood, Sean Bowe, Taylor Hornby, and Nathan Wilcox. Zcash protocol specification, version 2021.2.16 [nu5 proposal], 2021.
- Hop19. Daira Hopwood, 2019. <https://github.com/daira/tweedle>.
- Hop20. Daira Hopwood, 2020. <https://github.com/zcash/pasta>.
- KZG10. Aniket Kate, Gregory M. Zaverucha, and Ian Goldberg. Constant-size commitments to polynomials and their applications. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 177–194. Springer, Heidelberg, December 2010.
- KZM<sup>+</sup>15. Ahmed Kosba, Zhichao Zhao, Andrew Miller, Yi Qian, Hubert Chan, Charalampos Papamanthou, Rafael Pass, abhi shelat, and Elaine Shi. C0c0: A framework for building composable zero-knowledge proofs. Cryptology ePrint Archive, Report 2015/1093, 2015. <https://ia.cr/2015/1093>.
- Lee21. Jonathan Lee. Dory: Efficient, transparent arguments for generalised inner products and polynomial commitments. In *Theory of Cryptography Conference*, pages 1–34. Springer, 2021.
- Lip16. Helger Lipmaa. Prover-efficient commit-and-prove zero-knowledge SNARKs. In David Pointcheval, Abderrahmane Nitaj, and Tajjeeddine Rachidi, editors, *AFRICACRYPT 16*, volume 9646 of *LNCS*, pages 185–206. Springer, Heidelberg, April 2016.
- LRR<sup>+</sup>19. Russell W. F. Lai, Viktoria Ronge, Tim Ruffing, Dominique Schröder, Sri Aravinda Krishnan Thyagarajan, and Jiafan Wang. Omniring: Scaling private payments without trusted setup. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *ACM CCS 2019*, pages 31–48. ACM Press, November 2019.
- Poe18. Andrew Poelstra, 2018. <https://moderncrypto.org/mail-archive/curves/2018/000992.html>.
- WTs<sup>+</sup>18. Riad S. Wahby, Ioanna Tzialla, abhi shelat, Justin Thaler, and Michael Walfish. Doubly-efficient zkSNARKs without trusted setup. In *2018 IEEE Symposium on Security and Privacy*, pages 926–943. IEEE Computer Society Press, May 2018.
- ZBK<sup>+</sup>22. Arantxa Zapico, Vitalik Buterin, Dmitry Khovratovich, Mary Maller, Anca Nitulescu, and Mark Simkin. Caulk: Lookup arguments in sublinear time. Cryptology ePrint Archive, Paper 2022/621, 2022. <https://eprint.iacr.org/2022/621>.

## A Forest Single-Level Select-and-Rerandomize

Here we describe a slight generalization of the relation in Section 4.2.1. Its underlying relation takes as input hidden indexes  $j \in [\eta]$ ,  $i \in [\ell]$  and  $\eta$  public roots of a forest of  $(\ell, \mathbb{E}_1, \dots, \mathbb{E}_d)$ -Curve Trees  $C_1, \dots, C_\eta \in \mathbb{E}_d$ . In this context we also assume a root  $\hat{C}$  of a  $(\ell, \mathbb{E}_1, \dots, \mathbb{E}_{d-1})$ -Curve Tree which is a (known<sup>18</sup>) rerandomization of one of the children of  $C_1, \dots, C_\eta \in \mathbb{E}_d$ . Practically, this is accomplished by opening the  $j$ 'th commitment  $C_j$  to  $\vec{x}, \vec{y}$  using a Pedersen commit-and-prove over  $\mathbb{E}_d$ , then rerandomizes the  $i$ 'th point  $\hat{C} = (\mathbf{x}_i, \mathbf{y}_i) + [\delta] \cdot H^{(d-1)} \in \mathbb{E}_{d-1}$ . Slightly more formally we require a zero-knowledge argument of knowledge for the following relation:

$$\mathcal{R}^{(1\text{-Lvl-CrvT-Frst})} := \left\{ (i, j, r, \delta, \vec{x}, \vec{y}) : \begin{array}{l} C_j = \langle \vec{x}, X^{\vec{t}} \rangle + \langle \vec{y}, Y^{\vec{t}} \rangle + [r] \cdot H^{(d)} \\ \wedge \hat{C} = (\mathbf{x}_i, \mathbf{y}_i) + [\delta] \cdot H^{(d-1)} \end{array} \right\}$$

*Remark 3 (Simple  $\sqrt{n}$  Membership Proof).* Note that  $\mathcal{R}^{(1\text{-Lvl-CrvT-Frst})}$  immediately provides a very simple membership proof with  $O(\sqrt{n})$  prover/verifier complexity using  $\eta$   $(\ell, \mathbb{E}_1, \mathbb{E}_2)$ -Curve Trees and letting  $\eta = \ell = \sqrt{n}$ : to prove that a Pedersen commitment  $\hat{C} = C_q + [\delta] \cdot H^{(1)} \in \mathbb{E}_1$  is a rerandomization of a commitment from the list  $C_1^{(1)} = (\mathbf{x}_1^{(1)}, \mathbf{y}_1^{(1)}), \dots, C_n^{(1)} = (\mathbf{x}_n^{(1)}, \mathbf{y}_n^{(1)}) \in \mathbb{E}_1$ . Let  $\eta = \ell = \sqrt{n}$  and for each  $i \in [\eta]$  compute:  $C_i^{(2)} = \langle \vec{x}_{[(i-1) \cdot \ell : i \cdot \ell]}^{(1)}, X^{(1)} \rangle + \langle \vec{y}_{[(i-1) \cdot \ell : i \cdot \ell]}^{(1)}, Y^{(1)} \rangle$  i.e. commit to each chunk of  $\ell$  children individually. Since the size of the “select and rerandomize” relation above is  $O(\eta + \ell)$ , by applying a proof system with (quasi)-linear verification time we obtain a membership proof with  $\tilde{O}(\sqrt{n})$  complexity.

## B Circuit Specifications

*Remark 4 (Custom Gates).* We keep the explication of our techniques as broadly applicable as possible: working for any elliptic curve on short wierstrass form and any commit-and-proof system for Pedersen commitments. However, the circuits in this section can be further optimized for particular curves (e.g. with non-trivial efficient endomorphisms) and proof systems (e.g. Plonk [GWC19] with custom gates for elliptic curve operations, and/or, Plookup [GW20]).

<sup>18</sup> In terms of extraction.

We provide all circuit specifications as Rank-1 constraints systems (R1CS): the left side of any constraint consists of a product ( $\times$ ) of affine combinations, while the right side consists of an affine combination.

### B.1 2-Set Membership

To constrain  $w \in \{v_1, v_2\}$ , enforce the following R1CS constraint:

$$(w - v_1) \times (w - v_2) = 0 \quad (1)$$

Most commonly  $w \in \{0, 1\}$  (i.e.  $v_1 = 0$  and  $v_2 = 1$ ).

### B.2 Not Zero

To enforce  $v \neq 0$ , introduce  $t_1$  and constrain:

$$t_1 \times v = 1 \quad (2)$$

### B.3 Curve Check

For a point  $P = (x, y) \in \mathbb{E}(\mathbb{F})$ , introduce  $t_1, t_2$  and constraints:

$$x \times x = t_1 \quad (3)$$

$$x \times t_1 = t_2 \quad (4)$$

$$y \times y = t_2 + Ax + B \quad (5)$$

### B.4 Incomplete Curve Addition

We denote by  $\div$ : incomplete addition on the short Weierstrass curve  $\mathbb{E}$ , formally:

$$\mathbb{E} \cup \{\perp\} \div \mathbb{E} \cup \{\perp\} \rightarrow \mathbb{E} \cup \{\perp\}$$

$$\perp \div \_ \mapsto \perp$$

$$\_ \div \perp \mapsto \perp$$

$$1 \div \_ \mapsto \perp$$

$$\_ \div 1 \mapsto \perp$$

$$P \div -P \mapsto \perp, P \in \mathbb{E}$$

$$P \div P \mapsto \perp, P \in \mathbb{E}$$

$$P \div Q \mapsto P + Q, P \in \mathbb{E}, Q \in \mathbb{E}, \text{ Otherwise}$$

In other words: for points  $(x_1, y_1), (x_2, y_2) \in \mathbb{E}(\mathbb{F})$  the operation is undefined when  $x_1 = x_2$  (and undefined on points not on the curve) or when one of the operands is the point at infinity. For three points (witnesses)  $(x_1, y_1), (x_2, y_2), (x_3, y_3)$  we enforce  $(x_3, y_3) = (x_1, y_1) \div (x_2, y_2)$ , by introducing a free variable for the slope  $\delta$  and the 3 constraints:

$$\delta \times (x_2 - x_1) = y_2 - y_1 \quad (6)$$

$$\delta \times (x_3 - x_1) = -y_3 - y_1 \quad (7)$$

$$\delta \times \delta = x_3 + x_1 + x_2 \quad (8)$$

### B.5 Checked Curve Addition

When exceptional cases may occur, we can check for these by enforcing distinct  $x$ -coordinates. i.e. to enforce:

$$(x_3, y_3) = (x_1, y_1) + (x_2, y_2)$$

Enforce the constraints:

$$x_1 \neq x_2 \quad (9)$$

$$(x_3, y_3) = (x_1, y_1) \div (x_2, y_2) \quad (10)$$

## B.6 Secret 3-Bit Lookup

An  $n$ -dimensional secret lookup in a constant table, i.e.  $v = T[b_0 + 2 \cdot b_1 + 2^2 \cdot b_2]$  for secret  $b_0, b_1, b_2 \in \{0, 1\} \subseteq \mathbb{F}$  and  $v \in \mathbb{F}^n$  with  $T : \mathbb{N}_8 \rightarrow \mathbb{F}^n$ . For a table  $T : \mathbb{N}_8 \rightarrow \mathbb{F}$  the lookup requires 5 R1CS constants:

$$b_0 \in \{0, 1\} \tag{11}$$

$$b_1 \in \{0, 1\} \tag{12}$$

$$b_2 \in \{0, 1\} \tag{13}$$

$$b_{\&} = b_1 \times b_2 \tag{14}$$

$$b_0 \times \begin{pmatrix} -T_0 \cdot b_{\&} + T_0 \cdot b_2 + T_0 \cdot b_1 - T_0 + T_2 \cdot b_{\&} \\ -T_2 \cdot b_1 + T_4 \cdot b_{\&} - T_4 \cdot b_2 - T_6 \cdot b_{\&} \\ +T_1 \cdot b_{\&} - T_1 \cdot b_2 - T_1 \cdot b_1 + T_1 - T_3 \cdot s_{\&} \\ +T_3 \cdot b_1 - T_5 \cdot b_{\&} + T_5 \cdot b_2 + T_7 \cdot b_{\&} \end{pmatrix} = \begin{pmatrix} T_b - T_0 \cdot b_{\&} + T_0 \cdot T_2 + T_0 \cdot b_1 - T_0 + T_2 \cdot b_{\&} \\ -T_2 \cdot b_1 + T_4 \cdot b_{\&} - T_4 \cdot b_2 - T_6 \cdot b_{\&} \end{pmatrix}$$

In general, for tables  $T : \mathbb{N}_8 \rightarrow \mathbb{F}^n$  the technique above requires  $4 + n$  constants: repeating the last constraint for each additional coordinate.

## B.7 Circuit for Fixed-Base Exponentiation

Abusing notation, we write  $(\tilde{x}, \tilde{y}) = (\mathbf{x}, \mathbf{y}) \div T$  for the constraint:  $(\tilde{x}, \tilde{y}) = (\mathbf{x}, \mathbf{y}) \div (\hat{x}, \hat{y})$  and  $(\hat{x}, \hat{y}) \in T$ . Multiplying a constant curve point by a secret scalar is implemented by decomposing the scalar into 3-bit windows  $(b_0, b_1, b_2)$  and defining the tables  $T$  st. the exceptional cases does not occur (except for the last table – where we use the checked version). Let  $m = \lfloor \lambda/3 \rfloor + 1$ , for for  $i \in 1, \dots, m-1$ , define the table  $T_i$  as:

$$T_i = \left\{ \left[ j \cdot 2^{3 \cdot i} + 2^{3 \cdot (i+1)} \right] \cdot H \mid j \in 0, \dots, 2^3 - 1 \right\}$$

Define  $T_m$  as follows:

$$T_m = \left\{ \left[ j \cdot 2^{3 \cdot m} - \sum_{i=1}^{m-1} 2^{3 \cdot i} \right] \cdot H \mid j \in 0, \dots, 2^3 - 1 \right\}$$

To enforce  $(\tilde{x}, \tilde{y}) = [r] \cdot H + (\mathbf{x}, \mathbf{y})$ , we express it as:

$$(\tilde{x}, \tilde{y}) = \text{Rerand}(\mathbf{x}, \mathbf{y}) := (\tilde{x}, \tilde{y}) = (\mathbf{x}, \mathbf{y}) + (T_m + (T_{m-1} \div (T_{m-2} \div (\dots)))) \tag{15}$$

And decompose with witness  $r \in \mathbb{Z}_{|H|}$  as

$$r = \sum_i v_i \cdot 2^{3i}$$

## B.8 Range Check

A range check for  $v \in [0, 2^i)$  requires  $i$  constraints:

$$\forall b_i \in \{0, 1\} \tag{16}$$

$$v = \sum_i 2^i \cdot b_i \tag{17}$$

## B.9 Selection

Selecting a single secret entry (hidden index) from a secret vector:

$$\mathbf{x} = \text{Select}(\vec{\mathbb{X}}) := \{(\mathbf{x}, \vec{I}, \vec{\mathbb{X}}) : \forall j : I_j \in \{0, 1\} \wedge \mathbf{x} = \sum_{j=1} \mathbb{X}_i \cdot I_i \wedge 1 = \sum_{j=1} I_j\}$$



## B.10 Spending Relation

The spending relation consists of a relation for each curve in the tower of curves (see Section 3.1). Below we describe the spending relation for a two cycle of curves  $(\mathbb{E}, \hat{\mathbb{E}})$ . The spending relation for the commit-and-prove over  $\mathbb{E}$ :

$$\text{Spend}(v, \text{pk}, C_1) := \left\{ ((\vec{\mathbb{X}}_i, \hat{x}_i, \hat{y}_i)_{i \in 1, \dots, a/2}, v) : \right.$$

$$\begin{aligned} \text{Root} = C_1 = (\mathbb{x}_1, \mathbb{y}_1) &= \text{Commit}(\vec{\mathbb{X}}_1) && // \text{ Open commitment to vector of } \mathbb{x}\text{-coordinates} && (18) \\ \hat{x}_1 &= \text{Select}(\vec{\mathbb{X}}_1) && // \text{ Select a coordinate.} && (19) \\ (\hat{x}_1, \hat{y}_1) &\in \mathcal{P}_{\hat{\mathbb{E}}} && // \text{ Decompress to "permissible" point.} && (20) \\ \hat{C}_1 = (\hat{x}'_1, \hat{y}'_1) &= \text{Rerand}(\hat{x}_1, \hat{y}_1) && // \text{ Rerandomize inner commitment.} && (21) \\ &\vdots && && \\ C_{d/2-1} = (\mathbb{x}_{d/2-1}, \mathbb{y}_{d/2-1}) &= \text{Commit}(\vec{\mathbb{X}}_{d/2-1}) && (22) \\ \hat{x}_{d/2-1} &= \text{Select}(\vec{\mathbb{X}}_{d/2-1}) && (23) \\ (\hat{x}_{d/2-1}, \hat{y}_{d/2-1}) &\in \mathcal{P}_{\hat{\mathbb{E}}} && (24) \\ \hat{C}_{d/2-1} = (\hat{x}'_{d/2-1}, \hat{y}'_{d/2-1}) &= \text{Rerand}(\hat{x}_{d/2-1}, \hat{y}_{d/2-1}) && (25) \\ &\vdots && && \\ \hat{C}_{d/2} &= \text{Commit}(\text{pk}, v) && // \text{ Partially open the final commitment. } v \text{ is hidden.} && \left. \right\} && (26) \end{aligned}$$

Note that  $\text{pk}$  is a part of the statement (public input). The spending relation for the commit-and-prove over  $\hat{\mathbb{E}}$ :

$$\left\{ ((\vec{\mathbb{X}}_i, \mathbb{x}_i, \mathbb{y}_i)_{i \in 1, \dots, a/2}) : \right.$$

$$\begin{aligned} \hat{C}_1 = (\hat{x}_1, \hat{y}_1) &= \text{Commit}(\vec{\mathbb{X}}_1) && // \text{ Open commitment to vector of } \mathbb{x}\text{-coordinates} && (27) \\ \mathbb{x}_1 &= \text{Select}(\vec{\mathbb{X}}_1) && // \text{ Select a coordinate.} && (28) \\ (\mathbb{x}_1, \mathbb{y}_1) &\in \mathcal{P}_{\mathbb{E}} && // \text{ Decompress to "permissible" point.} && (29) \\ C_1 = (\mathbb{x}'_1, \mathbb{y}'_1) &= \text{Rerand}(\mathbb{x}_1, \mathbb{y}_1) && // \text{ Rerandomize inner commitment.} && (30) \\ &\vdots && && \\ \hat{C}_{d/2} = (\hat{x}_{d/2}, \hat{y}_{d/2}) &= \text{Commit}(\vec{\mathbb{X}}_{d/2}) && (31) \\ \mathbb{x}_{d/2} &= \text{Select}(\vec{\mathbb{X}}_{d/2}) && (32) \\ (\mathbb{x}_{d/2}, \mathbb{y}_{d/2}) &\in \mathcal{P}_{\mathbb{E}} && (33) \\ C_{d/2} = (\mathbb{x}'_{d/2}, \mathbb{y}'_{d/2}) &= \text{Rerand}(\mathbb{x}_{d/2}, \mathbb{y}_{d/2}) && \left. \right\} && (34) \end{aligned}$$

The statement (public input) is defined by  $x = (\text{pk}, (C_i, \hat{C}_i)_{i \in 1, \dots, a/2})$ . Note that  $C_1 = \text{Root}$  is the root of the Curve Tree containing all prior transaction outputs.

## B.11 Minting Relation

The minting relation is only over  $\mathbb{E}$  and consists of a range check:

$$\text{Mint}(v, \text{pk}, C) := \{(v, \text{pk}) : \quad (35)$$

$$C = \text{Commit}(v, \text{pk}) \quad (36)$$

$$v \in [0, 2^{64}] \quad (37)$$

Note that the  $\text{pk}$  is part of the witness.

## B.12 2-to-2 UTXO Relation

The 2-to-2 UTXO relation consists of two instances of the spending relation, two instances of the minting relation and a check that the balances (minus a transaction fee) match.

$$\text{UTXO}(\text{Root}, C_{\text{out1}}, C_{\text{out2}}) := \left\{ (v_{\text{in1}}, v_{\text{in2}}, v_{\text{out1}}, v_{\text{out2}}, \text{pk}_{\text{out1}}, \text{pk}_{\text{out2}}) : \right. \quad (38)$$

$$\text{Mint}(v_{\text{out1}}, \text{pk}_{\text{out1}}, C_{\text{out1}}) \quad (39)$$

$$\text{Mint}(v_{\text{out2}}, \text{pk}_{\text{out2}}, C_{\text{out2}}) \quad (40)$$

$$\text{Spend}(v_{\text{in1}}, \text{pk}_{\text{in1}}, \text{Root}) \quad (41)$$

$$\text{Spend}(v_{\text{in2}}, \text{pk}_{\text{in2}}, \text{Root}) \quad (42)$$

$$v_{\text{in1}} + v_{\text{in2}} = v_{\text{out1}} + v_{\text{out2}} + v_{\text{fee}} \quad (43)$$

Note that  $\text{pk}_{\text{in1}}$  and  $\text{pk}_{\text{in2}}$  are public, while  $\text{pk}_{\text{out1}}$  and  $\text{pk}_{\text{out2}}$  are part of the witness.